

Research Statement of Hamsa Balakrishnan

The air transportation system is a complex, global system that transports over 2.1 billion passengers each year. Air traffic delays have become a huge problem for passengers and airlines; they even make the headlines in the popular press and lead to new legislation. Aircraft are also the fastest growing contributor to man-made greenhouse gas emissions. According to the Joint Economic Committee of the US Senate, domestic air traffic delays in 2007 cost airlines over \$19 billion and the US economy over \$41 billion, wasted 740 million gallons of jet fuel, and released an additional 7.1 billion kilograms of CO₂ into the atmosphere.

My research is in the design, analysis, implementation, and evaluation of *practical algorithms for air transportation systems* to help air traffic controllers and system operators make better decisions in the face of increasing traffic. This research is important because of the high costs of delays and pollution today, as well as the projected doubling in air traffic over the next fifteen years. To prevent cascading delays and congestive collapse, and to mitigate pollution, we need new techniques and strategies, perhaps even a radical redesign of certain aspects of the system.

My research style is to develop new algorithms that are grounded in real-world data, implement them, and test them in both simulation and in field trials to gain a fundamental understanding of which techniques work well and why. By analyzing flight, weather, and operational datasets, I have (together with collaborators) developed algorithms for several tasks, including:

1. Scheduling and routing of aircraft in the air and on the ground, accounting for multiple objectives and stakeholders.
2. Coping with the intrinsic uncertainty in the operating conditions (e.g., weather).
3. Incorporating environmental objectives into decision-making.

The rest of this statement summarizes my contributions in these areas, my teaching activities, and future plans. References to papers are given in parentheses; detailed citations are listed in the complete list of publications.

Research activities and contributions

My research tackles air traffic management problems at all stages of flight, including on the airport surface, takeoffs and landings, and the airborne/en-route phase. It also addresses the system-level effects that arise from the interactions between the different elements in the system.

I. Airport Operations

Airport congestion management with Pushback Rate Control: Aircraft taxiing on the surface contribute significantly to the fuel burn and emissions at airports. My research (with my students Ioannis Simaiakis, Harshad Khadilkar and Melanie Sandberg, and John Hansman and Tom Reynolds) identifies opportunities to reduce airport congestion, designs and field-tests surface management strategies, and estimates the benefits of these strategies [*Transp. Res. Record* 2010].

The central idea is to control the rate at which aircraft push back from their gates at those times when the airport is congested. If better control were to be exercised, aircraft will spend time at their gates with their engines off, instead of adding to an already congested taxiway system with their engines on.

The question is how to make this intuition work. To solve this problem, Ioannis Simaiakis and I developed and validated a new stochastic queuing network model of departures. Given the aircraft pushback schedule, our model predicts the expected taxi-out time and queuing delay for

each flight, as well as the congestion levels, runway schedules, and departure throughput of the airport. This model includes an estimation of unimpeded taxi-out time distributions, and uses the transient analysis of $D(t)/E_k(t)/1$ queuing systems [AIAA-GNC 2009, AIAA-GNC 2010; Simaiakis PhD Thesis 2013].

We conducted field trials of Pushback Rate Control at Boston's Logan Airport (BOS) with the help of the Massachusetts Port Authority and the FAA. The results from two phases of field tests conducted in 2010-2011 showed that during fifteen four-hour demonstration periods, more than 23,000 kg of fuel were saved, at the rate of 50-60 kg per gate-held flight. Moreover, these savings were achieved with average gate-hold times of only 4.7 minutes [ATM-R&D-Seminar 2011 (best paper award), ICRA 2012, ACC 2012, *IEEE Transactions on Intelligent Transportation Systems* 2013]. In 2011, we tested a new variant of Pushback Rate Control using approximate dynamic programming, with an interface for air traffic controllers implemented on an Android tablet computer. This work won the inaugural CNA (formerly the Center for Naval Analysis) Award for Operational Analysis (2012), which "recognizes work that is judged as having provided the most creative, empirically based support to a real-world decision, or solution to a real-world problem."

The methodology developed for BOS has also been adapted to operations at other airports [ATM-R&D-Seminar 2013]. With the support of the FAA, we are currently preparing for a demonstration of these surface management approaches at LaGuardia Airport (LGA).

Metrics to characterize airport operational performance: Harshad Khadilkar and I have developed hybrid multi-modal estimation algorithms for processing airport surface surveillance data to determine aircraft ground trajectories, locations where aircraft queue up, the queue characteristics, and the resultant wait times. We have applied these algorithms to data from BOS, Dallas (DFW), New York's LaGuardia (LGA), and Philadelphia (PHL) airports. At BOS, we have used these analyses to develop tools that provide air traffic controllers feedback on their performance. Since 2011, we have provided these results to the BOS Operations Manager in the form of daily operational efficiency reports [AIAA-ATIO 2011, *ATC Quarterly* 2013].

Surface traffic management from a network control perspective: In recent work, Harshad Khadilkar and I showed how pushback control can be formulated as a network congestion control problem and solved efficiently using approximate dynamic programming. We have shown how this approach can effectively address practical resource constraints such as limited gate capacity [ACC 2012, *AIAA Journal of Guidance, Control and Dynamics* 2013, ECC 2013].

II. Arrival and Departure Operations

Practical multi-objective scheduling algorithms: Scheduling takeoffs and landings on runways is challenging because it needs to address three competing considerations: efficiency, safety, and equity among airlines. A natural approach to runway scheduling is Constrained Position Shifting (CPS) [Dear and Odoni], which requires that an aircraft's position in the scheduled sequence not deviate significantly from its position in the first-come-first-served sequence. With Bala Chandran and my student Hanbong Lee, I developed a new family of scalable dynamic programming algorithms for runway scheduling under CPS and other operational constraints [AIAA-GNC 2006; ATM-R&D-Seminar 2007; ACC 2008; *Operations Research* 2010]. The key insight is that even though the space of all feasible solutions is exponential in size, we can represent the solution space as a directed acyclic graph (DAG) whose size is linear in the number of aircraft being scheduled. This insight enabled us to reduce scheduling under CPS to a shortest path problem on that DAG.

We have developed a prototype implementation, which is fast enough for real-time use. We have also shown how this framework can be extended to many practical problems, including the simultaneous scheduling of takeoffs and landings, and the optimization of more general cost functions, including fuel burn and robustness metrics [ACC 2007; *Proc. of the IEEE* 2008]. Our implementation has been released to researchers at the NASA Ames Research Center, and has been integrated with NASA's Stochastic Terminal Area Scheduling Simulation for evaluating future operating concepts. Our methods have, for the first time, made scheduling under CPS a practical way to increase terminal-area throughput.

New mechanisms for Collaborative Decision Making: A Ground Delay Program (GDP) is initiated when congestion is expected at an airport to allocate the available airport capacity among the scheduled flights. In the first step of a GDP, a static or dynamic stochastic ground holding problem is solved in order to determine the ground delay and arrival slot assigned to each flight. The Collaborative Decision Making (CDM) framework then allows airlines to redistribute the slots assigned by ground-holding models to their flights, depending on flight-specific delay costs. My student Varun Ramanujam and I have identified a tradeoff between the ability of the ground holding model to dynamically adapt to forecast updates, and the flexibility to redistribute slots during the CDM step. As a result, an allocation that is optimal before the application of CDM mechanisms may be suboptimal afterwards. We proposed a new hybrid stochastic ground-holding model that combines the desirable properties of the static and dynamic models. Using a range of realistic case studies, we demonstrated that the hybrid stochastic ground-holding model yields a greater reduction in delay costs over a range of possible GDP scenarios [Ramanujam PhD Thesis 2011].

I have also studied market-based mechanisms for slot exchanges between airlines, and evaluated the nature of incentives for airlines to participate in these mechanisms and to report their true preferences, as well as the susceptibility of slot allocation mechanisms to manipulation by airlines [IEEE-CDC 2007].

III. Airspace operations

Robust routing of air traffic flows: Convective weather (thunderstorms) is responsible for large delays and disruptions in many parts of the world. Current flight scheduling and routing algorithms require reliable weather forecasts. My student Diana Michalek Pfeil and I showed how to translate raw convective weather forecasts, which provide deterministic predictions of the Vertically Integrated Liquid (VIL, the moisture content of a region of airspace), into probabilistic forecasts of whether or not a route into or out of an airport will be blocked.

Meteorologists predict the VIL on a scale of 0-256 around the airport, for each pixel on a 1 km x 1 km grid. Valid routes must remain on pixels with VIL less than 133. An aircraft that is asked to fly a route that ends up being blocked by weather will have to be diverted significantly and rescheduled. Predicting whether a route will be open for use is tricky because the VIL forecasts are inaccurate, and routes are much longer in duration than the granularity of VIL forecasts.

Using techniques from machine learning, we developed and validated classification algorithms that predict whether or not a given route is likely to be open in actual weather [AMS-Annual-Meeting 2009, ATM-R&D-Seminar 2009]. Our approach uses historical forecasts and the characteristic features of the route. Ours is the first algorithm that combines different features of the route to predict the probability of blockage, and provides several insights into the relationship between VIL forecasts and route blockage. Surprisingly, we found that the theoretical capacity (a measure of how many routes into the airport do not pass through forecast weather obstacles) was

a poor predictor of route blockage, despite being a frequently cited metric. The reason is that although the theoretical capacity is a prediction of how many routes will be open, it does not give any indication of which ones they would be.

We have also used our forecasts of route blockage to modify routes dynamically to optimize the expected capacity of the terminal-area. Experiments using real weather scenarios show that our algorithms recommend that a terminal-area route be modified 30% of the time, opening up 11% of available routes that would have otherwise been closed. We also found that 97% of routes predicted by our method as being “open with probability greater than 95%” are in fact open in the weather that actually materializes [CDC 2010, *Transportation Science* 2011].

A key implication of our work is that improved metrics that assess the skill in predicting route blockage are a better measure than traditional metrics, which focus on the accuracy of predicting convective activity in a 1 sq km pixel. We are currently identifying factors that drive pilot behavior, for example, why pilots sometimes fly through “Level 5” weather, while others deviate from “Level 2” weather. Factors such as delays, airlines, and demand are considered in this work, which uses a combination of flight trajectory data and weather archives.

IV. System-level challenges

Prediction of air traffic delays: My student Juan Jose Rebollo and I have developed a new air traffic delay prediction model that incorporates both temporal (time-of-day, day-of-week, etc.) and network delay states (the overall condition of the National Airspace System or NAS) as explanatory variables. We used clustering to determine six “typical delay states.” These delay states are intuitive, corresponding to times when delays are high in the New York, Chicago, or Atlanta areas. We examined the prevalence of certain types of delay days during certain months of the year and used this information in our prediction algorithms. For the 100 most delayed origin-destination (OD) pairs in the NAS, the average error in predicting (two hours in advance) whether the mean departure delay on that link will exceed 60 minutes was only 19%. In addition, the average test error increased by only 3.5% for a prediction horizon of 6 hours [ICRAT 2012].

Distributed feedback control of the NAS: Today’s airspace is partitioned into sectors, and each air traffic controller is responsible for managing traffic within his/her sector. Controllers only communicate locally with their neighboring sectors; the control of flows between sectors is done through an ad hoc negotiation of handoffs. Because flows are not prioritized, local weather disruptions in one area (say, New York) can lead to holding patterns far away (the mid-West), impacting all flows in the area, even those not bound for New York.

With Jerome Le Ny, I developed a queuing network representation of traffic flows that accurately models current operations. We developed distributed feedback control techniques that guarantee that the aircraft queues in each airspace sector, which are an indicator of controller workload, are kept small. We showed that under realistic conditions, our feedback control policy for scheduling and routing aircraft stabilizes the system (all queues remain bounded). Our approach provides the first distributed feedback control strategy for realistic, multi-airport settings. We also showed how our methods could be used to mitigate the impact of weather disruptions [ACC 2009, CDC 2010, *AIAA Journal of Guidance, Control and Dynamics* 2011].

Combined communication and control algorithms for air traffic management: In recent work with Harshad Khadilkar, Pangun Park, and Claire Tomlin, I have addressed the problem of designing combined communication and control protocols for air traffic control. The research seeks to find the level of decentralization that balances system safety and efficiency. For example, ADS-B surveillance can potentially be used to shift air traffic control to a more

distributed architecture; however, channel variations and interference with existing secondary radar replies can affect ADS-B systems. We design and simulate a protocol that combines centralized control in congested regions with distributed control in low traffic regions, and show that its performance is comparable to fully centralized strategies, despite requiring much less ground infrastructure [*IEEE Transactions on Intelligent Transportation Systems*, 2013].

Future Plans

In the long term, I am interested in the development of practical and principled approaches to design robust and sustainable infrastructure systems.

Many existing infrastructures (e.g., transportation, energy, communications, etc.) are large-scale, multi-stakeholder systems, and face similar challenges with regards to day-to-day operations and proposed improvements. Realistic models of these systems are essential for the design of implementable algorithms to improve their performance. The approaches that I have developed for air transportation, where I use diverse operational data sets to develop appropriate models and algorithms for control and optimization, can be extended to other infrastructures as well. In addition to developing new control strategies, these methodologies can be used to evaluate past performance [*Air Traffic Control Quarterly* 2013, *Transportation Research Part D*, AIAA-ATIO 2013] and also predict future behavior [ICRAT 2012].

Infrastructure systems typically involve complex interactions among their different components; I am therefore interested in investigating the interactions between different parts of the system, the integration of scheduling algorithms for different elements, and the design of architectures that improve overall system performance [*IEEE Trans. on Intelligent Transportation Sys.* 2013].

The interaction between strategic market-based approaches to resource allocation (such as, landing slot auctions) and more tactical control strategies (such as, congestion management) is another exciting topic that I am investigating. These problems have traditionally been considered independently, even though the strategic allocations influence the tactical control strategies, and vice versa. An understanding of these interactions will not only lead to better air traffic management algorithms, but also have parallels in other domains (such as electricity markets).

Finally, a key element of many infrastructures is the presence of the human stakeholder. Machine learning approaches applied to data can help us understand the factors that influence the actions and decisions of humans in the system. Our initial studies in aviation, considering the decision processes of air traffic controllers and pilots, have yielded promising results [ACC 2010, *Transportation Research Record* 2013 (submitted)]. A deeper understanding of current decision making processes will ultimately help in the design of better decision support systems.

In summary, I believe that by using approaches from data mining, control systems engineering, optimization and game theory, we will be able to design, implement, and evaluate practical algorithms for a more efficient, robust, and green air transportation system.